reviewer's guide

HEWLETT PACKARD

HP ProCurve Hubs & Switches

HP ProCurve Switch Routing Switch 9304M
HP ProCurve Switch Routing Switch 9308M

Date 4/1/1999
Version 4.0

HP ProCurve Routing Switch 9308M

HP ProCurve Routing Switch 9304M

Less Work, More Network

Table of Contents

1. BACKGROUND 4

1.1 HP's Proactive Networking 4

1.2 What is a Routing Switch? 4

1.3 Hewlett-Packard Switch Products 4
1.3.1 HP ProCurve Routing Switches Covered in this Guide 4
1.3.2 Other Switches Available From HP 5

1.4 HP Switch Positioning 7
1.4.1 Positioning for the HP ProCurve Routing Switch 9304M 7
1.4.2 Positioning for the HP ProCurve Routing Switch 9308M 7

2. EVALUATION FEATURES AND BENEFITS 8

2.1 Feature Set Summary 8

2.2 Architecture 9
2.2.1 Hardware Architecture Summary 9
2.2.2 Packet Processor 9
2.2.3 Shared Memory 9
2.2.4 Forwarding Engine 9
2.2.5 Priority Queues 10
2.2.6 System Management Interface 10
2.2.7 Backplane Design 10

2.3 Layer 3 Services 10
2.3.1 IP 11
2.3.2 IPX 11
2.3.3 AppleTalk 11

2.4 QoS 12
2.4.1 IEEE 802.3x Flow Control 12
2.4.2 IEEE 802.1p Priority Support 12
2.4.3 Multicast Support 12
2.4.3.1 DVMRP 12
2.4.3.2 PIM (Dense) 12
2.4.3.3 IGMP 13
2.4.3.4 IP Tunneling 13
2.4.4 Layer 4 Features 13
2.4.4.1 Priority Queues 13
2.4.4.2 Layer 4 Filtering 13

2.5 VLANs 13
2.5.1 IEEE 802.1Q VLAN Support 14
2.5.2 IEEE 802.1D Spanning Tree Protocol 14

2.6 Port Trunking (Link Aggregation) 14

2.7 Filtering 15
2.7.1 Layer 2 Filtering 15
2.7.2 Layer 3 and Layer 4 Filtering 16

2.8 Flexibility **16**

2.9 Reliability **16**
2.9.1 Lifetime Software Updates (Best in the Industry) 17
2.9.2 Warranty - One Year 17
2.9.3 Availability 17
2.9.4 Service and Support 17
2.9.5 Year 2000 Compliance 18

2.10 Network Management **18**
2.10.1 MIB Support 18
2.10.2 RMON Support 19
2.10.3 Network Monitoring Port 19
2.10.4 Console Support 19

2.11 Some 9304M / 9308M Routing Switch Maximums **20**

3. PERFORMANCE TESTING **21**

3.1 Throughput/Packet Loss Rate **21**

3.2 Congestion Control **22**
3.2.1 Head-of-line-blocking 22
3.2.2 Bi-directional Congestion Testing (X-Stream) 22

3.3 Latency **22**

4. ADDITIONAL INFORMATION **24**

4.1 ProCurve Networking Web Site **24**

4.2 Reseller Plaza **24**

4.3 White papers **24**

5. PRICING **25**

6. APPENDIX A: PERFORMANCE FIGURES **26**

6.1 HP ProCurve Routing Switch 9308M **26**
6.1.1 IP Routing Throughput Test – Packets Within Modules 26
6.1.1.1 Gigabit-SX ports 26
6.1.1.2 100BT 26
6.1.2 IP Routing Throughput Test – All Packets Across the Backplane 26
6.1.2.1 Gigabit-SX ports 26
6.1.2.2 100BT 26
6.1.3 Latency Test - Unicast Traffic – Packets Within the Modules 27
6.1.3.1 Gigabit-SX ports 27
6.1.3.2 100Mbps 27
6.1.4 Latency Test – Unicast Traffic – Packets Across Backplane 27
6.1.4.1 Gigabit-SX ports 27
6.1.4.2 100Mbps[5] 27

1. Background

This Routing Switch Reviewer's Guide will help network testing engineers at computer trade publications, resellers and end-user sites evaluate the merits of the HP ProCurve routing switches.

1.1 HP's Proactive Networking

HP proactive networking is an extensible line of products working together to provide the control network administrators need to deliver the network uptime and performance that their organizations require. This guide describes one part of HP proactive networking: the HP ProCurve Routing Switch 9308M and HP ProCurve Routing Switch 9304M products.

1.2 What is a Routing Switch?

A routing switch combines the speed and low cost of a Layer 2 switch, with the ability to route at Layer 3. Unlike traditional routers which are processor based, a routing switch does its routing in hardware at or near wire speed for IP. Some routing switches, such as the HP ProCurve routing switches, also route IPX in hardware. In addition, HP ProCurve routing switches provide processor-based AppleTalk routing. Each port on the routing switch can be routed, or combinations of ports can be switched, with the port groups tied to virtual router interfaces inside the routing switch.

It is this combination of high speed, flexibility, and low cost that make routing switches very attractive for local LAN environments.

1.3 Hewlett-Packard Switch Products

1.3.1 HP ProCurve Routing Switches Covered in this Guide

This guide covers the following Hewlett-Packard routing switch products:

HP ProCurve Routing Switch 9308M

A feature-rich, modular, 8 slot chassis-based routing switch delivering 96 million pps performance with up to 64 Gigabit Ethernet ports or up to 188 10/100 Mbps autosensing ports. Routes IP, IPX and AppleTalk. This product is ideal for large networks and provides high performance as a collapsed backbone. Modules available are:

- HP ProCurve 9300 10/100 Module (24 port)
- HP ProCurve 9300 10/100 Management Module (16 port)
- HP ProCurve 9300 100Base-FX Module (24 port, MT-RJ connectors)
- HP ProCurve 9300 Gigabit SX Module (8 port, SC connectors)
- HP ProCurve 9300 Gigabit SX Management Module (8 port, SC connectors)
- HP ProCurve 9300 Gigabit SX/LX Module(4 ports SX, 4 ports LX, SC connectors)
- HP ProCurve 9300 Gigabit SX/LX Management Module (4 ports SX, 4 ports LX, SC connectors)

HP ProCurve Routing Switch 9304M

A feature-rich, modular, 4 slot chassis-based routing switch delivering 48 million pps performance with up to 32 Gigabit Ethernet ports or up to 88 10/100 autosensing ports. This product is ideal for medium-to-large networks and provides high performance in collapsed backbones, data centers, and server farms. Has the same architecture design and uses the same modules as the 9308M.

Pricing for each of these switches and modules is given in section 5.

1.3.2 Other Switches Available From HP

Below is a complete list of switches available from HP. Separate Reviewer's Guides are available for the HP ProCurve Switch 8000M / 1600M, HP ProCurve Switch 4000M / 2424M, and the HP AdvanceStack Switch 2000 / 800T switches.

All managed HP switches include embedded Web-based management and the HP TopTools for Hubs & Switches network management application.

HP ProCurve Switch 8000M

A modular 10 slot mini-chassis backbone Layer 2 switch that provides high 10/100 autosensing port density (up to 80), and also provides multiple Gigabit and 100Base-FX fiber-optic port connectivity. Advanced software features include Switch Meshing and Cisco Fast EtherChannel[1] support. Excellent as a small backbone switch. Modules available are:

- HP ProCurve Switch 10/100Base-T module - (8) autosensing 10/100 Ethernet UTP ports
- HP ProCurve Switch Gigabit-SX module - (1) 1000Mbps multi-mode fiber optic port
- HP Switch Gigabit-LX module provides (1) 1000Mbps multi-mode or single-mode fiber-optic port. Fully adheres to the IEEE 802.3z specification.
- HP ProCurve Switch 100Base-FX module - (4) 100Mbps multi-mode fiber optic ports
- HP ProCurve Switch 10Base-FL module - (4) 10Mbps multi-mode fiber optic ports

HP ProCurve Switch 1600M

A fixed port switch with sixteen fixed 10/100Mbps autosensing ports and one expansion slot, primarily for Gigabit connectivity, but will accept any of the modules listed with the HP ProCurve Switch 8000M above. Same software feature set as the Switch 8000M. Ideal for server farm connectivity.

HP ProCurve Switch 4000M

The HP ProCurve Switch 4000M is a modular switch that comes standard with (40) 10/100 Base-T UTP autosensing ports, with five open slots for additional modules. The HP ProCurve Switch can support up to (80) 10/100 autosensing 10/100Base-T ports, uses the same modules as the Switch 8000M, and is ideal for scaleable/expandable low-cost migration to 10/100 desktop switching for smaller groups of users.

HP ProCurve Switch 2424M

The HP ProCurve Switch 2424M provides 24 fixed 10/100 Base-T autosensing ports and one expansion slot that accepts any of the 8000M/1600M/4000M modules. There is also a HP ProCurve Switch 2424M Stacking Module available that provides two full-duplex transceiver-based Gigabit ports. Transceivers available are Gigabit SX and LX and a Stacking Kit that provides a low cost stacking connection between HP ProCurve Switch 2424Ms. The HP ProCurve Switch 2424M is ideal for low cost migration 10/100 desktop switching for smaller groups of users.

HP ProCurve Switch 2224

Fixed configuration unmanaged desktop switch with (24) autosensing 10/100 Mbps ports, one of the UTP ports can be used as a 100Base-FX port with the addition of an optional transceiver.

HP ProCurve Switch 212M

Fixed configuration top-of-stack switch with (12) 10Mbps Ethernet ports and (2) 100Base-T ports (one fixed TX and one transceiver-based port)

[1] EtherChannel® is a registered trademark of Cisco Systems, Inc.

HP ProCurve Switch 224M

Fixed configuration desktop switch with (24) 10Mbps Ethernet ports and (2) 100Base-T ports (one fixed TX and one transceiver-based port)

HP AdvanceStack Switch 800T

Eight fixed autosensing half/full duplex 10/100Mbps Ethernet ports, all transceiver-based for top-of-stack applications.

HP AdvanceStack Switch 2000

A flexible six slot high speed backplane switch for connectivity into legacy LAN environments. Modules available include:

- Two port 100Base-T module (transceiver-based),
- Two port 100VG module (transceiver-based),
- Four port 10Base-T module,
- Four port 10Base-FL module
- One port DAS FDDI module, and
- 1 port OC-3 ATM module (co-developed with FORE Systems).

1.4 HP Switch Positioning

HP identifies three different classifications for switches: node/desktop, top-of-stack and backbone switches. The routing switches are targeted at the medium to large backbone. Requirements in this area are high bandwidth, Layer 3 routing services, VLAN services, quality of service (QoS) controls, flexible filtering, and fault tolerance. The HP ProCurve Routing Switch 9304M and HP ProCurve Routing Switch 9308M meet all these needs, providing a stable, resilient, non-blocking backbone.

The HP ProCurve routing switches are well complemented by the Layer 2 HP ProCurve switches that do the wiring closet, server farm, and desktop aggregation, uplinking to the routing switch through either Gigabit Ethernet or 100 Mbps links or trunks.

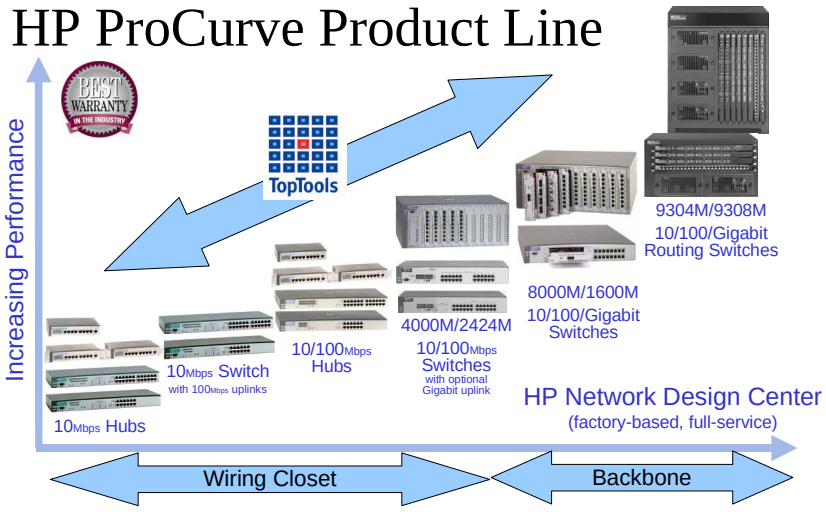

1.4.1 Positioning for the HP ProCurve Routing Switch 9304M

The HP ProCurve Routing Switch 9304M is a feature-rich, modular, chassis-based Layer 3 routing switch that delivers 48 million packets per second performance for up to 32 Gigabit ports or 88 10/100 Mbps ports. Layer 4 controls are also available, allowing tuning of packet streams based on Layer 4 information without needing the cooperation of the end nodes. Port flexibility is provided through the modular design, with modules available for 10/100 UTP autosensing connections, Gigabit SX and LX connections, and 100Base-FX connections. Designed for medium-to-large networks, the HP ProCurve Routing Switch 9304M provides non-blocking high performance Layer 3 routing in collapsed backbones, data centers and server farms.

1.4.2 Positioning for the HP ProCurve Routing Switch 9308M

The HP ProCurve Routing Switch 9308M is a feature-rich, modular, chassis-based Layer 3 routing switch that delivers 96 million packets per second performance on up to 64 Gigabit ports or 188 10/100 Mbps ports. Layer 4 controls are also available, allowing tuning of packet streams based on Layer 4 information without needing the cooperation of the end nodes. Port flexibility is provided through the modular design, with modules available for 10/100 UTP autosensing connections, Gigabit SX and LX connections, and 100Base-FX connections. Designed for large networks, the HP ProCurve Routing Switch 9308M provides non-blocking high performance Layer 3 routing at the core of the network.

2. Evaluation Features and Benefits

2.1 Feature Set Summary

The HP ProCurve Routing Switch 9304M and HP ProCurve Routing Switch 9308M are store-and-forward Layer 3 routing switches that feature:

- A non-blocking Layer 3 architecture consisting of shared memory on each card, all interconnected via the backplane cross point matrix switch fabric. 9308M: 96 Million pps, 9304M: 48 Million pps.
- Hardware wire speed routing for IP and IPX. Processor-based 300,000 pps routing for AppleTalk.
- IP Functionality supported:
  - Routing services: RIP (v1, v1 compatible v2, and v2), OSPF (RFC 1583 and RFC 2178 compliant).
  - DNS Resolver - converts names to DNS server addresses
  - UDP Helper - supports packet relay of UDP packets for certain services
- IPX RIP/SAP routing supported
- AppleTalk Routing Table Maintenance Protocol (RTMP) routing
- Multicast support: IGMP, PIM (dense) and DVMRP
- Link Aggregation (Port Trunks) – Switch-to-switch and switch-to-server aggregated links allow scalable bandwidth communication. Can be used in many cases to trunk to non-HP devices. Compatible with Cisco Fast EtherChannel (PAgP turned off).
- Up to (4096) 802.1Q compliant port-based, Layer 2 and Layer 3 VLANs for segmentation of local LANs
- Extensive filtering capabilities based on Layer 2 addresses and Layer 3 routes or addresses. For IP, Layer 4 port number filtering is supported. AppleTalk zone filtering is available. Filtering (except AppleTalk) is performed at wire-speed.
- Virtual Router Redundancy Protocol (VRRP) – Provides hot standby router services. A second 9300 can be in active stand-by on the network in case of failure of the primary router. VRRP is a proposed standard of the IETF, RFC 2338.
- One year warranty, support services are available for additional years.
- Free lifetime software updates (best in the industry)
- Free end-user telephone support available worldwide during the warranty period.
- Extensive net management capabilities including:
  - Web-based management for anytime, anywhere configuration access
  - HP TopTools for Hubs & Switches (standalone), seamless integration with HP OpenView/NT
  - Integration into HP OpenView/UX with option HP Hub & Switch Management for HP OpenView/UX.
  - SNMP/MIB II/RMON support on all ports for monitoring and control
  - Ability to configure a network monitoring port for use with external probes or analyzers
- Spanning Tree Protocol (802.1D) support - one instance per VLAN.
- Year 2000 compliant

2.2 Architecture

2.2.1 Hardware Architecture Summary

The HP ProCurve Routing Switch 9304M has 4 identical slots, while the HP ProCurve Routing Switch 9308M has eight. Any of the modules can be put in any of the slots. The only rule is that there must be one, and only one, card with the management function installed in the chassis. While there are managed versions of each of the different modules available (except for the 100Base-FX module), it is best to have management on a Gigabit card as there is no loss of actual ports. The managed version of the 10/100 card, on the other hand, has only 16 ports on it compared to the 24 ports on the unmanaged card.

Each of the modules, which is based on a shared memory structure, is interconnected through a backplane using a crosspoint matrix interconnect. The crosspoint matrix in the HP ProCurve Routing Switch 9308M is twice as large as the HP ProCurve Routing Switch 9304M to handle up to double the number of packets expected from the eight slots versus the four slots in the 9304M.

The HP ProCurve Routing Switch 9304M has two slots for the load-sharing power supplies. One supply ships with the 9304M and can power a fully loaded chassis. The HP ProCurve Routing Switch 9308M has four power supply slots available; a minimum of two supplies (supplied) are needed to run a fully loaded chassis. A third and/or fourth power supply can be installed for redundancy and longer expected overall power supply life.

Overall, the hardware is designed to provide a non-blocking architecture, so that the HP ProCurve routing switches can easily sit in the middle of a medium-to-large enterprise campus providing the traffic-forwarding performance needed.

The HP ProCurve routing switches allow 64,000 MAC addresses. Boot time from a cold start is <10 seconds.

2.2.2 Packet Processor

Each port on a module has a packet processor associated with it. The packet processor is responsible for reading into the packet headers so that forwarding decisions can be made based on the Layer 2, Layer 3, Layer 4 and filtering requirements. The packet processor then generates a forwarding identifier (FID) that defines destination port, port mirroring requirements, packet type, VLAN affiliation, prioritization and other parameters. If processor based functions are needed, such as AppleTalk routing, that is also determined. The packet processor also modifies the bits in the packet as required for routing, such as time-to-live (TTL), MAC addresses and checksums. If Layer 4 information is to be processed, that is also determined at this point, with the FID modified as necessary.

Once the FID has been determined for a packet, it is stored in a cache area that is managed by the management board. Actual movement of the packet through the routing switch data pathways is specified by the FID, resulting in the packet making it's way to the proper output port(s) of the routing switch.

2.2.3 Shared Memory

Modules are based on a shared memory architecture. Packets are buffered on the module in this shared memory area. The pipeline to/from the packet processors to the shared memory has a 64 Gbps throughput. The actual data rate through this switch fabric is 32 Gbps, since in shared memory designs the packet has to travel the fabric twice, once inbound and once outbound. The 32 Gbps is ample, as the maximum data rate the shared memory will see is 16 Gbps, 8 Gbps from the ports, and 8 Gbps from the backplane. The packet is managed in shared memory through a shared memory identifier (SMID).

This shared memory area size on the module is 2 to 8 MB, depending on the module type.

2.2.4 Forwarding Engine

The forwarding engine, a custom ASIC present on each module, uses the information in the FID to actually move the packets to and from either external ports or the backplane switch fabric via the shared memory. The forwarding engine also determines the priority queue (described in the next section) designation for outbound packets. The forwarding engine has its own separate data path to the packet processors to make packet processing time as short as possible.

2.2.5 Priority Queues

Priority queues form part of the QoS implementation in the routing switch. Each external port on a module has 4 priority queues that define the order outbound packets will be sent to that port. There are also 4 priority queues on the module for each of the other modules installed in the chassis. These are used for packets destined to travel the backplane to another module. The priority queues themselves only contain the SMIDs as pointers to the actual packets in shared memory.

These priority queues map into the priority designation of 802.1p. Since 802.1p has 8 levels of priority possible, 802.1p priority levels 0 and 1 map to the module priority of 0 (lowest), 802.1p levels 2 and 3 to module priority 1, and so on. Packets without any 802.1p tagging are normally assigned module priority 0, unless the port priority factors have been modified through user configuration. The priority queues are managed through a fair-weighted queuing algorithm that prevents any priority queue from getting starved due to high traffic levels in other queues.

Through user configuration, priority of packets can be changed based on the packet's MAC address, VLAN affiliation, port the packet came in on, IP address, or IP port number.

2.2.6 System Management Interface

The System Management Interface, present on the management card, stores and maintains the port and system-level master configuration tables, routing tables, Layer 2, Layer 3, and Layer 4 address tables, and all FID registers. Packet functions requiring processor attention, such as AppleTalk routing or broadcast packet handling are also handled through this interface.

Communication from the individual module packet processors to the System Management Interface is done through the management bus, a control path distinct from the packet data path. Any modification of the FID for those packets requiring additional central processor attention are also handled through the management bus. The central processor runs at 240 MHz.

The System Management Interface has 32 MB of DRAM memory available. This memory is used for switch tables, such as routing tables, ARP cache and MIB tables. At factory default, many of the table sizes are less than their maximum in order to conserve memory. Tables can be resized through user configuration, if necessary. Default and maximum values are given in this guide for some tables.

2.2.7 Backplane Design

The modules are interconnected through a backplane crosspoint switch matrix, allowing packets a direct path to their destination module. Each module in the chassis has an 8 Gbps path in both directions to this matrix.

The advantage of a crosspoint matrix is that all modules can send packets to every other module simultaneously. The crosspoint matrix can handle about 22 Mpps in the 9304M and 44 Mpps in the 9308M.

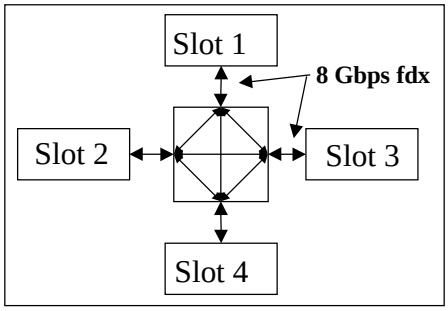

2.3 Layer 3 Services

Routing for both IP and IPX is done in ASIC hardware at wire speed. AppleTalk routing is done by the processor at a rate up to 300,000 pps. Routing services can be applied two different ways:

- The router is directly assigned to specific ports – these ports look like traditional router ports, where each port is in its own IP subnet or IPX network
- A group of ports are assigned to a VLAN and tied to a virtual router interface inside the routing switch – externally these ports are switched together (since they are in the same VLAN) and are all in the same IP subnet or IPX network. Internal to the routing switch they are tied to the routing service through the virtual router port without the loss of an external port. Groups of these VLANs can communicate with each other through their virtual internal router interfaces. Sixty virtual router interfaces are allowed in the 9304M, 180 in the 9308M.

  The use of router virtual interfaces provides a high degree of flexibility with the routing switch and is called Integrated Switch Routing (ISR).

2.3.1 IP

Since there is no performance degradation, IP routing is, by default, enabled on the HP ProCurve routing switches. Some of the IP services available are:

- Routing Services
  - RIP (version 1, version 1 compatible version 2, and version 2)
    - Split Horizon and Poison Reverse supported
    - Redistribution (importing of OSPF or static routes into the RIP route table)
  - OSPF (RFC 1583 and 2178 compliant)
    - Load sharing of equal cost paths
    - Redistribution (importing of RIP or static routes into the OSPF route tables)
    - OSPF traps (RFC 1850)
  - Static IP routes (64 maximum per chassis)
  - IRDP (ICMP Router Discovery Protocol)
- Proxy ARP and RARP
- Bootp Relay Service
- UDP broadcast forwarding (UDP Helper)
- Encapsulation type: Ethernet II (default) or SNAP
- DNS Resolver (allows users to use host name rather than IP address for services such as Telnet, ping or trace route.)
- Extensive IP-based filtering available (see the Filtering section)

2.3.2 IPX

- Routing services: RIP/SAP
- Four encapsulation types: Ethernet SNAP, Ethernet 802.2, Ethernet 802.3 and Ethernet II
- NetBIOS broadcast routing support
- IPX filters (see the Filtering section)

2.3.3 AppleTalk

The HP ProCurve routing switches support AppleTalk Phase II. AppleTalk is processor-based in the 9304M and 9308M. Performance is up to 300,000 packets per second. If multiple ports are being switched together and then routed for AppleTalk through a virtual interface to these switched ports, the 300,000 pps rate is for traffic through the virtual interfaces into the router. Traffic across the

switched ports is at wire speed, since it is looked at as Layer 2 traffic that just happens to be AppleTalk at Layer 3.

- Routing Services
  - Routing Table Maintenance Protocol (RTMP)
  - AppleTalk Echo Protocol (AEP)
  - Name Binding Protocol (NBP)
- Datagram Delivery Protocol (DDP)
- AppleTalk Address Resolution Protocol (AARP)
- Zone Information Protocol (ZIP)
- AppleTalk-based filters (see the Filter section below)

2.4 QoS

2.4.1 IEEE 802.3x Flow Control

Both HP ProCurve routing switches adhere to the IEEE 802.3x Flow Control specification. This provides industry standard support for managing heavy port traffic and helps the routing switch avoid memory buffer overflows. Flow control is enabled on a port-by-port basis. By default, all ports have flow control enabled on the HP ProCurve routing switches.

2.4.2 IEEE 802.1p Priority Support

IEEE 802.1p packet tagging supports both designation of VLAN membership (see the VLAN section below) and packet priority (up to 8 levels). The architecture of the HP ProCurve routing switches support four levels of priority through multiple priority queues each module keeps. See the priority queue discussion in the previous section for more details on the handling and mapping of the priority queues. These priority queues allow the 9304M / 9308M to be responsive to time-sensitive applications that use the priority field in packet tagging.

Priority of packets can be given specific values based on their inbound port, VLAN affiliation or MAC address. This allows the network manager to better manage the traffic through the switch, giving priority to packets that contain time-sensitive data without involvement of the end nodes.

The 9300 routing switches can work in concert with the other HP managed switches to provide an end-to-end QoS solution. The HP desktop and local backbone switches with firmware code C.07.01[2] or later can use the IP address, VLAN affiliation, or IP/TOS (type of service) field of the packet to set a value for the 802.1q priority bits. As noted above, the 9300 routing switches can read these bits and prioritize the packets appropriately through the routing switch. Through the combination of the HP desktop and local backbone switches, and the 9300 routing switches, packet priorities can be established end-to-end in the local LAN environment without any end node interaction.

2.4.3 Multicast Support

The routing switches support three areas of multicast packet management, DVMRP, PIM and IGMP. DVMRP and PIM are routing protocols used to set up the route trees through the network for multicast groups when multiple routes exist. IGMP is used where there are not multiple paths, such as at Layer 2 or in a single or dual router environment, and can work in conjunction with either DVMRP or PIM, since these protocols will establish a single path for any one multicast stream. DVMRP and PIM, each tied to different ports, can run concurrently in the same routing switch.

2.4.3.1 DVMRP

Distance Vector Multicast Routing Protocol (DVMRP) is a protocol similar to RIP in that builds its own multicast routing table in conjunction with other routers in the network. DVMRP uses this table to

[2] Available June 1, 1999. Available for download at no cost to present switch owners.

generate multicast path trees through the network so that a particular multicast stream is only sent on links that ultimately reach a receiver of that stream. This prevents multicast streams from being sent to parts of the network where it isn't needed. DVMRP builds and maintains a separate multicast tree for each of the offered multicast streams.

2.4.3.2 PIM (Dense)

Protocol-Independent Multicast (PIM) is very similar to DVMRP except that it uses the IP routing table of the switch rather than build its own table. Since it doesn't build a table directly, it doesn't need to send out routing protocol packets to the other routers in the network like DVMRP does, hence the protocol-independent name. The HP ProCurve routing switches support the dense form of PIM, the version useful in LAN applications where there may be a high concentration of receiving nodes.

PIM generates multicast path trees to transmit a multicast stream only on those links leaving the routing switch that have a receiver for that multicast stream somewhere downstream.

2.4.3.3 IGMP

Internet Group Multicast Protocol (IGMP) is another multicast control protocol that builds delivery trees, but does not keep or use any routing information in this process, making it useful for Layer 2 environments. IGMP, handling the Layer 2 environment, is used in conjunction with either DVMRP or PIM, which handle the Layer 3 environment. Both routing switches support Version 2 of IGMP and can act as an IGMP querier to determine who belongs to the multicast groups, as do all of HP's managed switches. The routing switch can then direct a specific multicast stream to only those segments that have nodes that have joined the multicast group associated with that stream.

The HP ProCurve routing switches support 256 IGMP groups at default, 1024 with a user configurator change.

2.4.3.4 IP Tunneling

Not all routers support DVMRP or PIM. If there are a series of non-supporting routers in-between the HP ProCurve routing switches, IP tunneling can be used for multicast traffic. This will prevent the multicast traffic from being treated as broadcast traffic on the non-supporting routers. (It will also prevent the any potential multicast clients connected to these intervening routers from seeing the multicast streams.) This will allow the DVMRP routing functions to work correctly, since logically the two routing switches will appear to be neighbors, even though physically they may be separated by a number of intervening routers.

2.4.4 Layer 4 Features

2.4.4.1 Priority Queues

As mentioned in section 2.2.5, each port on the 9300 routing switches has four priority queues. While these priority queues are automatically used for packets having an 802.1q priority field, the 9300 routing switches can also be configured by the user to specify the use of a particular priority queue based on the packet's IP address or IP port number.

This capability provides QoS for a packet based on its Layer 4 information. Applications using particular port numbers can be given priority treatment through the switch without the end node clients having to be aware of QoS, particularly valuable since client operating systems and the applications themselves are generally not QoS aware at the moment.

Layer 4 control of priority is a very powerful feature for applications that use isochronous (time-sensitive) data streams, such as Voice-over-IP or conferencing. This feature can also be used to give high priority to multicast applications.

Other applications can be prioritized as appropriate. For example, http traffic could be prioritized lower through the switch, or data involved in backups prioritized higher.

2.4.4.2 Layer 4 Filtering

The 9300 routing switches can filter based on IP addresses, IP port number and IPX network number. More details are given later in the Filtering section. These filters can be used in permit or deny situations, allowing fine tuning of traffic based on Layer 4 information.

2.5 VLANs

A Virtual LAN is a logical collection of ports or nodes that belong to a single broadcast/multicast domain. VLANs were originally devised as a solution to limit the size of any one broadcast domain to allow scaling of switched environments. With the advent of routing switch solutions, however, use of VLANs in end user environments is now largely done for network policy or security reasons. VLANs are also used in the HP ProCurve routing switches to establish groups of ports that are switched, linked to the router through a router virtual interface. See the Layer 3 services above for a more detailed explanation.

The HP ProCurve routing switches support up to 4096 VLANs (8 default), although normal VLAN usage is usually less than twenty VLANs. VLAN membership can be designated through any one of the following:

- A particular port (port-based)
- A 802.Q tag
- IP protocol
- IP subnet
- NetBIOS
- IPX network number
- AppleTalk
- Decnet
- Other

VLANs can overlap on a single port. For example, it may be advantageous to have a server connected through a single port to be a member of two different VLANs such that two different groups of PCs can access the same server, but the two groups of PCs cannot talk directly with each other.

Port-based VLANs can be further subdivided by using protocol VLANs. Protocol VLANs establish packet membership based on the packet's IP subnet number, IPX network number, etc. Since a packet could be part of several VLANs simultaneously there is a hierarchy of VLAN ownership. Port-based VLANs are the lowest level. Layer 3 protocol-based VLANs, IP, IPX, AppleTalk, DECnet and NetBIOS are in the middle. IP sub-net and IPX network number VLANs are at the top.

VLANs can also be assigned to the virtual interfaces of the router in the routing switch. This provides a means of communication between two VLANs. VLANs defined only on switched ports within the routing switch cannot talk to each other without going through the router.

2.5.1 IEEE 802.1Q VLAN Support

The HP ProCurve routing switches support the IEEE 802.1Q VLAN tagging standard. The routing switches can have multiple VLAN traffic share a single physical link. 802.1Q also allows interoperability at this level between different vendors in a standards-based way.

Ports with only a single VLAN designation can be designated as untagged ports. Packets leaving these ports will not be 802.1Q tagged.

2.5.2 IEEE 802.1D Spanning Tree Protocol

The HP ProCurve routing switches support multiple instance spanning tree protocol. Each VLAN can have an instance of spanning tree running. This is useful in multiple VLAN environments where loops exist for redundancy purposes at the Layer 2 level.

Note that other HP switches support a single instance of spanning tree per box, per the 802.1D spec. If an HP Layer 2 switch is connected to a HP ProCurve routing switch and both are running spanning tree, the 9300 port connected to the HP switch needs to be an untagged port for the spanning tree protocols to be recognized, so that within the routing switch VLAN a port will be behave correctly according to Spanning Tree.

2.6 Port Trunking (Link Aggregation)

Link Aggregation is the industry term for the ability to combine multiple coterminous links (links that begin at the same point and end at the same point) as one logical link.

Link aggregation allows two HP ProCurve routing switches to be interconnected by 2-to-4 of the same type of links, with all links acting as one higher-speed link. Since the number of links in a trunk is configurable, the bandwidth is scaleable to the needs of a particular network. For example, (4) 100Mbps links can be trunked to provide the equivalent of a 400 Mbps (800 Mbps full-duplex) link between two switches or 4 Gigabit links can be trunked for the equivalent of a 4 Gigabit (8 Gigabit full-duplex) link. Fiber-optic links can be trunked to interconnect switches across large campuses. Port trunking also provides redundancy on links between the two switches or switch and server. If one of the links fails, the traffic is moved to another link in the trunk in under one second.

The 9304M / 9308M have two forms of trunking, switch and server. Switch trunking balances traffic across the links based on destination address, server trunking uses the source addresses for load balancing. The MAC address is used if the trunked ports on the routing switch are switched, IP or IPX addresses are used for routed ports. Use the switch form of trunking to interconnect switches and the server version of trunking to interconnect the routing switch with a server.

The HP ProCurve Routing Switch 9304M and HP ProCurve Routing Switch 9308M have been certified for use with HP-UX server trunked ports.

Up to 4 links can be used in a trunk. All links in a particular trunk must be on the same module, contiguous, start with a port on a MOD4 boundary (1,5,9,13, 17, or 21) and have the same characteristics, i.e. duplex, QoS, etc. Based on these requirements, the 9304M can have a maximum of 22 trunk groups, the 9308M can have 46 trunk groups.

Trunks on the routing switches are compatible with the HP Trunks on other HP ProCurve switches; use the DA (destination address) setting. These trunks are also compatible with Cisco Fast EtherChannel trunks; use the server form of the trunk on the 9304M / 9308M and turn PAgP off on the FEC side.

2.7 Filtering

Both the HP ProCurve Routing Switch 9304M and HP ProCurve Routing Switch 9308M support a variety of Layer 2 and Layer 3 filters. With the exception of AppleTalk and broadcast packets, filtering is done in hardware so there is no performance penalty for using filters. With filter hierarchy defined by the user, filtering is very flexible.

All filters can be specified for individual ports, groups of ports, or globally. This flexibility allows definition of very specific rules for how a packet should be handled.

Filter maximums: 1024 per chassis, 128 on any one port.

2.7.1 Layer 2 Filtering

Filters at Layer 2 are:

- Inbound and/or outbound MAC addresses
- Frame types

The frame type field can be used to filter on Layer 3 protocol types without having the router functionality of the routing switch involved.

2.7.2 Layer 3 and Layer 4 Filtering

At Layer 3 and Layer 4, these filters available are:

IP Filters

- Inbound and/or outbound IP addresses
- Protocol types (UDP, TCP, etc.)
- Port numbers or ranges of numbers (ftp, http, etc.)
- Specific RIP routes inbound
- Specific advertised RIP routes outbound
- RIP redistribution (conversion of OSPF routes to RIP)
- OSPF redistribution (conversion of RIP routes to OSPF)

IPX Filters

- Inbound and/or outbound IPX addresses
- Inbound and outbound forward filters (Allows a remote IPX client access to a restricted-access server)
- Specific RIP routes inbound
- Specific advertised RIP routes outbound
- Specific advertised SAP routes (128 filters max per chassis)

AppleTalk Filters

- Zone filters - ability to filter resource access from other zones
- Network filters - Specific advertised zones in the routing packets

2.8 Flexibility

The HP ProCurve routing switches have been designed with flexibility and high port density in mind. Any module type can be plugged into any of the module slots. The only requirement is that one, and only one, module with management must be installed. The modules are hot swappable. Available module types are:

- HP ProCurve 9300 10/100 Module (24 port)
- HP ProCurve 9300 10/100 Management Module (16 port)
- HP ProCurve 9300 100Base-FX Module (24 port, MT-RJ connectors) MT-RJ is a new connector type that allows full duplex 100Base-FX connections in the same space as a RJ-45 connection. Adapter cables for SC to MT-RJ conversion are available from 3rd party sources.
- HP ProCurve 9300 Gigabit SX Module (8 port, SC connectors)
- HP ProCurve 9300 Gigabit SX Management Module (8 port, SC connectors) Works with 62.5/125 multi-mode fiber cable at up to 220M, or 50/125 multi-mode fiber cable at up to 500m.
- HP ProCurve 9300 Gigabit SX/LX Module(4 ports SX, 4 ports LX, SC connectors) SX connections work with 62.5/125 multi-mode fiber cable at up to 220M, or 50/125 multi-mode fiber cable at up to 500m connections. LX connections work with single-mode fiber cable up to 5km, or with either 62.5/125 or 50/125 multi-mode fiber cable up to 550m (a mode-conditioning cable may be required for multi-mode use on some older fiber-optic cables per IEEE 802.3z).
- HP ProCurve 9300 Gigabit SX/LX Management Module (4 ports SX, 4 ports LX, SC connectors) Same comments as for the card above.

A fully loaded HP ProCurve Routing Switch 9304M can have (88) 10/100Mbps ports or (32) Gigabit ports in a box only 22.8cm (9") in height and 13.8kg (30.8lbs).

A fully loaded HP ProCurve Routing Switch 9308M can have (188) 10/100Mbps ports or (64) Gigabit ports in a box only 52.7cm (20.75") in height and 26kg (57.3lbs).

2.9 Reliability

Overall reliability of a product is determined not only by how much time between product failures, but also how fast a failed unit can be repaired or replaced. HP offers a one year warranty, free basic-level phone support during this time, and lifetime free software updates. Some details are provided in this section.

2.9.1 Lifetime Software Updates (Best in the Industry)

Both the HP ProCurve Routing Switch 9304M and HP ProCurve Routing Switch 9308M come with lifetime software updates. For as long as software updates are available for these routing switches they can be downloaded from the HP web site for free. This includes not only bug fixes, but new software capabilities as well. This is an industry-leading product feature.

2.9.2 Warranty - One Year

Warranties, and the ease of obtaining warranty service for the end user, is a product benefit that is easily overlooked in a technical evaluation, but ranks high as a concern of end users as they get ready to actually make a purchase decision. The HP ProCurve Routing Switch 9304M and HP ProCurve Routing Switch 9308M have a one year warranty. If any part of the switch fails, including the power supply or fans, it will be replaced. In most parts of the world, the replacement unit is sent for next business day delivery in advance of the failing unit being returned to HP. Advance replacement gets the unit to the end user as fast as possible and minimizes down time by allowing the impaired unit to continue to be used if possible until the replacement unit arrives. This also allows for easy scheduling for when the actual unit swap occurs on the network.

The 9304m / 9308M one year warranty for a routing switch is industry-leading, with most other vendors offering 90 days.

Hewlett-Packard has support programs available to provide fee-based warranty-like service once the routing switch is out of warranty. See the Support and Service section for more details.

2.9.3 Availability

Availability is the measure of the ability for the switch to remain running over a period of time with minimal impact to the network environment.

Both the 9304M / 9308M allow hot-swapping of the port modules while the switch is still running without affecting the other port modules. This allows a defective port module to be replaced without affecting the rest of the network.

Optional redundant power supplies (identical to the primary supply(s)) can be installed. The 9304M allows one extra power supply. The 9308M allows up to four supplies to be installed, with at least two supplies at a minimum required. All power supplies load-share, allowing both supplies to run cooler, extending their Mean Time Between Failure (MTBF) values. If any of the supplies in a chassis with extra supplies (2 in the 9304M, 3 or 4 in the 9308M) does fail, the others continue to run preventing switch interruption. Power supplies can be hot-swapped (unplug it first!) during switch operation when extra supplies are present. Each supply has its own power cord.

The chassis fans are also hot-swappable.

2.9.4 Service and Support

HP provides free pre-sales and in-warranty post-sales installation and basic configuration telephone support during normal business hours to end users and HP resellers through the HP Customer Care Centers located world-wide.

In addition to free support services such as the warranty and telephone support, HP offers an extensive range of fee-based support services to meet more specialized needs. With the likelihood of the routing switches being installed at the core of the network, an on-site service upgrade is highly recommended.

On-site upgrades to either business-day four hour support or 24X7 four hour support can be purchased through HP.

Other fee-based services available that are highly recommended are hardware support contracts to cover the routing switches when they are no longer in warranty, and extensions to the basic phone support to cover network troubleshooting either during regular business hours or 24X7.

Hewlett-Packard can also provide site surveys, installation services, and actual management of the network, depending on customer needs.

More information can be found at http://www.hp.com/rnd/support/support.htm or by contacting the local HP sales office.

2.9.5 Year 2000 Compliance

The HP ProCurve Routing Switch 9304M and HP ProCurve Routing Switch 9308M are both Year 2000 compliant[3]. Refer to http://www.hp.com/year2000/compliance.html for HP's Year 2000 compliance definitions.

2.10 Network Management

Network management is an important part of a network solution. There are three levels of net management available for the HP ProCurve Routing Switch 9304M and HP ProCurve Routing Switch 9308M:

- Web-based management - Configuration of individual routing switches can be done anytime, anywhere through the web server available in each routing switch, accessible via a standard web browser.
- HP TopTools for Hubs & Switches - Management of a network of HP hubs, switches and routing switches can be done through the included HP TopTools for Hubs & Switches application, part of the HP TopTools suite of management products. HP TopTools for Hubs & Switches provides a network map and device configuration in an web-based format for the routing switches. HP TopTools for Hubs & Switches can manage up to 1500 nodes as shipped. The extra-cost HP TopTools Value Pak can be added that increases the number of monitored nodes to 5000.
- HP OpenView platforms - If management of a generic IP network is needed (i.e. management of HP and other vendors' devices):

  NT platform - bridge software to HP OpenView/NT is available so that TopTools can run under HP OpenView/NT. For more details see the TopTools web site at http://www.hp.com/toptools.

  UX platform - Management of the HP ProCurve Routing Switch 9304M and 9308M under HP OpenView/UX is available via the separately purchased J3250M HP Hub & Switch Management For HP OpenView-UX product.

TopTools will also run under Tivoli Enterprise and Tivoli NetView, as well as Unicenter TNG using the proper bridges to these environments. See the TopTools web site at http://www.hp.com/toptools for more details.

SNMP version 2 is supported.

[3] HP has made every effort to ensure the accuracy of our product testing. However, because each customer's environment is different from HP's laboratory test environment, it is the customer's responsibility to validate the Year 2000 readiness of these products in their own environment. Therefore, information about the Year 2000 status of HP products is provided "as is" without warranties of any kind and is subject to change without notice. HP makes no representation or warranty respecting the accuracy or reliability of information about non-HP products. Such information, if any, was provided by the manufacturers of those products and customers are urged to contact the manufacturer directly to verify Year 2000 readiness. The information provided here constitutes a Year 2000 Readiness Disclosure for purposes of the Year 2000 Information and Readiness Disclosure Act.

2.10.1 MIB Support

The HP ProCurve Routing Switch 9304M/ 9308M supports the following standard MIBs:

- MIB-II (RFC 1213)
- Ether-like MIB (RFC 1398)
- Bridge MIB (RFC 1493)
- Ethernet-like MIB (RFC 1643)
- RMON MIB (RFC 1757)
  Groups: (1) Ethernet statistics, (2) Ethernet history, (3) Alarm, (9) Event,

In addition, a number of enterprise-specific MIBs are also supported for such things as VLANs, and multiple bridge groups.

2.10.2 RMON Support

For those customers that use RMON applications, the HP ProCurve Routing Switch 9304M / 9308M support RMON groups 1 (Ethernet statistics), 2 (Ethernet history), 3 (Alarm), and 9 (Event). These four groups are available for all ports.

The Ethernet statistics group provides counters for packet counts, broadcast/multicast packets, packet length counts, missed packets and erred packets. Event and Alarm groups allow threshold setting and alarm generation based on the counters in the Ethernet group. History accumulates two records for each port, one a sampling of Ethernet statistics taken every 30 seconds and the other a sampling of statistics taken every 30 minutes.

2.10.3 Network Monitoring Port

If more RMON groups (such as packet capture) are desired, a RMON probe can be attached to one of the HP ProCurve Routing Switch 9304M / 9308M ports and Port Monitoring can be configured into the routing switch. Port Monitoring allows the end user to copy all traffic (inbound, outbound or both) from any number of ports (even those on different router ports) to a single destination port. This allows the probe to see all traffic on the selected ports to provide the probe with the proper global perspective.

LAN analyzers can also use this feature when doing specific monitoring or troubleshooting of network segments.

2.10.4 Console Support

Out-of-band management of HP routing switches can be done through the RS-232 console port via a directly connected terminal emulator. The console interface provides a command line interface to all configuration parameters.

The console interface is also available in-band through the telnet service.

The RS-232 cable necessary to connect with the 9304M / 9308M is a "straight-through" type and is supplied with the routing switch. It is a different cable than that used with other HP switches.

The RS-232 port is not speed-sensing. Default configuration is set at 9600 baud.

2.11 Some 9304M / 9308M Routing Switch Maximums

Parameter	Maximum Value	Default Value
MAC addresses	32,000	16,000
Port VLANs	4096	8
All filters	System: 1024 Interface: 128	
Layer 3 VLANs	1,024	32
Layer 4 sessions	4096	512
Virtual router interfaces	9304M: 60 9308M: 180	
ARP entries	16,000	4,000
IP cache size	64,000	16,000
IP routes	220,000	OSFP: 16,000 RIP: 40,000
IP route filters	System: 1,024 Interface: 128	System: 64 Interface: 32
Static routes	256	16
IGMP groups	1,024	255
DVMRP routes	64,000	2,048
PIM routes	64,000	16,000
IPX/SAP entries	8,192	4,096
IPX/RIP entries	8,192	2,048
IPX/SAP filters	128	32
IPX/RIP filters	128	32
IPX forward filters	128	32
AppleTalk routes	3,072	512
AppleTalk zones	System: 1024 Interface: 255	System: 256 Interface: 64
Number of Trunks	9304M: 22 9308M: 46	

3. Performance Testing

Routing switches are normally performance tested under a variety of conditions, at Layer 2 and Layer 3, for three main parameters: throughput/packet loss rate, congestion control and latency. Definitions for throughput, packet loss rate and latency are specified in RFC 1242. See http://www.rfc-editor.org/. Unfortunately, there is no clear definition for congestion control. There are numerous tests that test different aspects of congestion control, none of which is considered the definitive test for congestion.

Keep in mind that most of these tests are performed in an artificial environment intended to measure the outer limits of routing switches and rarely, if ever, reflect circumstances a switch would consistently find in an actual network environment. This is particularly true for congestion type testing. The applicability of these performance numbers to real world networks has to be interpreted carefully, as most routing switches in these tests perform to a level significantly exceeding the needs of real world networks.

Additional comments on each of the parameters follows.

3.1 Throughput/Packet Loss Rate

One of the major advantages that ASIC-based routing switches have over the traditional processor-based routers is that routing can be done near or at wire speed on all ports simultaneously for a cost considerably less than with a traditional router.

The throughput numbers given in Appendix A are shown for Gigabit-SX and 100Base-TX. The first table shows the situation where the packets don't have to leave the module they came in on. Since each module in the 9300 routing switches routes on the module itself, this first table shows throughput without any packets traveling the backplane. As shown in the table, the 9308M produces complete wire speed routing on all ports simultaneously. The maximum number of packets per second that can travel through the routing switch is also shown in the 64 byte packet size column at 95,238,080. Although not explicitly measured, this figure would be 47,619,040 for the 9304M (one-half the rate of the 9308M.) This figure was measured with the routing switch routing IP. It would be the same for IPX routing or just switching at Layer 2.

The numbers for 100Base-T are full media-speed as well, but since the maximum number of packets that can be introduced to the switch is limited by the total number of ports possible at 18.8 Gbps (27,976,280 pps), the routing engines are not taxed at all in this situation.

In the situation where the throughput test is set up such that the destination for all packets requires a trip across the crosspoint matrix backplane, throughput is at wire speed, except for 64 byte packets where throughput is 91.3% of wire speed. This is due to the crosspoint matrix handling of 86,956,544 pps maximum. Note that this is a number of packets distinction, not data throughput. At larger packet sizes where the actual number of packets is less, the data rate through the backplane is at full wire speed. While this is a limitation in the total throughput of the routing switch, this only happens when all packets cross the backplane, all packets are 64 bytes in length, and all ports are receiving packets at wire speed full duplex simultaneously. This situation only happens in a testing environment. (If it happens momentarily in a real network, the excess packets would be buffered in the modules' shared memory until the condition let up or shared memory became full.) To put this into perspective, if the average packet length through the routing switch exceeds 88 bytes, or if only 6 of the 8 gigabit ports on the modules have packets that cross the backplane, the backplane will run at full media speed. This clearly is not a limitation when dealing with real network environments.

This limitation is exactly the same for the 9304M. Since its crosspoint matrix is one-half the size of the 9308M, the matrix can handle 91.3% of media speed at 64 byte packets. The same comments given above as to why this really isn't a problem in real environments also holds for the 9304M.

For 100 Mbps environments, the crosspoint matrix backplane can handle the traffic at full wire speed. In this situation the traffic is limited by the total throughput of the 100 Mbps ports.

3.2 Congestion Control

Congestion occurs any time more packets are destined for a particular segment than that segment can handle. All switches handle momentary congestion passively through buffering excess packets in memory. The HP ProCurve Routing Switch 9304M and HP ProCurve Routing Switch 9308M handle the congestion issue largely through running at full media speed. The only congestion the routing switches need handle are packets destined for output ports that are congested due to external network segment congestion. For this momentary situation, shared memory on the modules is used. Each module has between 2 and 8 MB of shared memory, depending on the port type of the module.

Another standardized method to control congestion in a multiple network box environment is with the 802.3x Flow Control standard. The HP ProCurve routing switches support the 802.3x standard. Orderly control of segment traffic occurs because a port on the routing switch can signal its congestion to the segment, holding off segment traffic. For 802.3u to work correctly, the node(s) connected directly to routing switch ports with flow control enabled also have to support the flow control protocol. Flow control in the routing switches would only have to be used if the network segments external to the routing switches were congested.

In the past, some switches have implemented a non-standard form of congestion control called backpressure. The 802.3x Flow Control standard produces the same result as backpressure, but does it in a standardized way. HP switches do not implement backpressure.

3.2.1 Head-of-line-blocking

A particular type of congestion that can be seen in some network devices is head-of-line-blocking (HOLB). If a device exhibits HOLB, congestion on one outgoing port can affect the throughput of another outgoing port that is not congested. Whether a device exhibits HOLB or not is dependent on the inbound and outbound port buffer memory handling design. The shared memory architecture with fair-weighted queues on the HP ProCurve routing switches does not allow HOLB to occur.

3.2.2 Bi-directional Congestion Testing (X-Stream)

Another congestion test currently being used is the bi-directional, multi-stream spray test. In this test bursts of packet streams are sent into all tested ports at the same time. Collisions and congestion result simultaneously on all ports. Throughput and packet loss rate (PLR) are recorded for each port.

While this is an ardent stress test, the HP ProCurve routing switches handle it with little problem because of the wire-speed handling capability of the 9300 routing switches. Since X-Stream sends packets to all other ports except the port the packet came in on, not all packets travel the backplane.

3.3 Latency

Latency is the amount of time the switch takes to process a packet. In store-and-forward devices, such as with the HP ProCurve Routing Switch 9308M and 9304M, latency is reported by vendors on a LIFO basis (last bit in, first bit out), i.e. the time elapsed between the end of the last bit of the packet going into the switch to the beginning of the first bit of that packet emerging from the switch. Cut-through switch latency is measured on a FIFO basis (first in, first out), time measured from the beginning of the first bit of the packet going in to the switch to the beginning of the first bit of the outbound packet coming out of the switch. Latency figures should be used with care when comparing store-and-forward devices versus cut-through devices.

To add to the confusion, many testing houses report store-and-forward latency as FIFO because the test equipment reports it that way. Following the actual definition of latency (RFC1242), which is the time taken by the device to decide where to send a particular packet, store-and-forward routing switch

latencies should be reported as LIFO. At any rate, latency figures should be labeled as to whether they are LIFO or FIFO.

Occasionally store-and-forward routing switch latency will be reported on a FIFO basis. To convert FIFO to LIFO results, the time associated with the packet itself must be subtracted out. For 10Mbps Ethernet frames the following formula can be used:

LIFO = FIFO - (packet length (bytes) * .8 µsec)

For 100Mbps Ethernet the formula is:

LIFO = FIFO - (packet-length (bytes) * .08 µsec)

For Gigabit Ethernet the formula is:

LIFO = FIFO - (packet-length (bytes) * .008 µsec)

Latency is usually measured with the NetCom SmartBits or Ixia 1600 testers, special-purpose test units that can measure latency down to the nearest bit time for the technology under test.

In normal testing, latency should be measured for a variety of packet sizes and all ports should be tested since latency, being sensitive to the architecture of the switch, can vary from port to port.

As show by the figures in Appendix A, latency for the HP ProCurve Routing Switch 9308M and 9304M is excellent at around 7-17 µsec for packets forwarded on the same module and 10.5-21 µsec for packets that flow across the backplane. The difference in the 'on module' latency versus the 'across the backplane' latency is the time it takes to actually traverse the backplane and get stored on the target output module. Either way, these latencies are outstanding (very low). *And these latencies are with routing active.* The architecture of the HP ProCurve routing switches lead to the same packet latencies whether the routing switch is routing IP or IPX or Layer 2 switching the packets because these functions are done in the module ASICs. Multicast packets are handled in hardware incurring the same latency as unicast packets.

For those packets handled by the processor, such as AppleTalk packets or broadcast packets, the latency will run about 190-200 µsec, still a very respectable figure, particularly when compared to the much longer latencies usually seen for all packets in the traditional processor-based router.

4. Additional Information

4.1 ProCurve Networking Web Site

Additional information, including the latest data sheets, design guides, white papers, product documentation and support information can be obtained through the HP ProCurve Networking web site. HP ProCurve Networking can be reached at:

http://www.hp.com/go/procurve

4.2 Reseller Plaza

The dedicated reseller web site for HP Networking products can be reached either through the ProCurve Networking site, or directly at:

http://www.hp.com/go/sellprocurve

Admission to the reseller site requires registration.

4.3 White papers

White papers covering some of the technical details described here in more detail can be found in the Technical Library section of ProCurve Networking at:

http://www.hp.com/rnd/techlib/techlib.htm

5. Pricing

All HP routing switches ship with HP TopTools for Hubs & Switches.

Prod No.	Description	US List Price March 1, 1999
J4139A	**HP ProCurve Routing Switch 9304M chassis** Includes (4) open module slots	$7,999
J4138A	**HP ProCurve Routing Switch 9308M chassis** Includes (8) open module slot	$15,999
J4140A	HP ProCurve 9300 10/100 Module (24) autosensing UTP ports	$12,999
J4141A	HP ProCurve 9300 10/100 Management Module (16) autosensing UTP ports	$10,999
J4142A	HP ProCurve 9300 100Base-FX Module Includes (24) 100Base-FX ports, MT-RJ connectors	$24,999
J4143A	HP ProCurve 9300 Gigabit SX Module (8) multi-mode fiber ports, SC connectors	$25,999
J4144A	HP ProCurve 9300 Gigabit SX Management Module (8) multi-mode fiber ports, SC connectors	$26,999
J4145A	HP ProCurve 9300 Gigabit LX/SX Module (8) multi-mode fiber ports, SC connectors	$28,999
J4146A	HP ProCurve 9300 Gigabit LX/SX Management Module (8) multi-mode fiber ports, SC connectors	$29,999
J4147A	HP ProCurve 9300 Redundant Power Supply	$2,999
D7118A	**HP TopTools Value Pak** Allows TopTools management of 5000 nodes, up from standard 1500 nodes	$999

6. Appendix A: Performance Figures

These numbers have been generated by Hewlett-Packard, using two Ixia 1600 testers from Ixia Communications. Ixia testers are used by a number of network testing houses and the press to determine performance numbers for networking equipment. In these tests, 64 ports were used for Gigabit testing, 188 ports for 100 Mb testing.. All ports were full duplex. Numbers presented here are condensed from Ixia reports in order to save space.

Testing done on the HP ProCurve Routing Switch 9308M. Percentage of throughput would be the same for the 9304M, but at one-half the number of packets since the 9304M has on-half the number of ports of the 9308M.

See Section 3 for an analysis of these numbers.

6.1 HP ProCurve Routing Switch 9308M

6.1.1 IP Routing Throughput Test – Packets Within Modules

6.1.1.1 Gigabit-SX ports

Port pairs active, full duplex: 64 = 64 Gbps data out of the tester

Test length: 10 seconds

Packet size (bytes)	64	128	256	512	1024	1280	1518
Total Packets / sec	95,238,080	54,054,080	28,985,536	15,037,568	7,662,848	6,153,856	5,201536
Percentage	100.0	100.0	100.0	100.0	100.0	100.0	100.0

6.1.1.2 100BT

Port pairs active, full duplex: 188 = 18.8 Gbps data out of the tester

Test length: 10 seconds

Packet size (bytes)	64	128	256	512	1024	1280	1518
Total Packets / sec	27,976,280	15,878,480	8,514,520	4,417,436	2,250924	1,807,808	1,528,064
Percentage	100.0	100.0	100.0	100.0	100.0	100.0	100.0

6.1.2 IP Routing Throughput Test – All Packets Across the Backplane

6.1.2.1 Gigabit-SX ports

Port pairs active, full duplex: 64 = 64 Gbps data out of the tester

Test length: 10 seconds

Packet size (bytes)	64	128	256	512	1024	1280	1518
Total Packets / sec	86,956,544	54,054,080	28,985,536	15,037,568	7,662,848	6,153,856	5,201536
Percentage	91.3	100.0	100.0	100.0	100.0	100.0	100.0

6.1.2.2 100BT

Port pairs active, full duplex: 188 = 18.8 Gbps data out of the tester

Test length: 10 seconds

Packet size (bytes)	64	128	256	512	1024	1280	1518
Total Packets / sec	27,976,280	15,878,480	8,514,520	4,417,436	2,250924	1,807,808	1,528,064
Percentage	100.0	100.0	100.0	100.0	100.0	100.0	100.0

6.1.3 Latency Test - Unicast Traffic – Packets Within the Modules

6.1.3.1 Gigabit-SX ports

Port pairs active: 64
Latency test length: 10 seconds
Port-Pair Latency in microsec (µs)

Packet size (bytes)	64	128	256	512	1024	1280	1518	
Average of all ports	7.01	7.58	9.63	13.58	21.55	25.52	29.32	FIFO
	6.50	6.56	7.58	9.48	13.36	15.28	17.18	LIFO (Calculated)[4]

6.1.3.2 100Mbps[5]

Port pairs active: 188
Latency test length: 10 seconds
Port-Pair Latency in microsec (µs)

Packet size (bytes)	64	128	256	512	1024	1280	1518	
Port 1 to 2								FIFO
								LIFO (Calculated)[4]

6.1.4 Latency Test – Unicast Traffic – Packets Across Backplane

6.1.4.1 Gigabit-SX ports

Port pairs active: 64
Latency test length: 10 seconds
Port-Pair Latency in microsec (µs)

Packet size (bytes)	64	128	256	512	1024	1280	1518	
Average of all ports		12.06	11.59	16.06	25.42	29.52	33.91	FIFO
		11.04	9.54	11.96	17.23	19.28	21.76	LIFO (Calculated)[4]

6.1.4.2 100Mbps[5]

Port pairs active: 188
Latency test length: 10 seconds
Port-Pair Latency in microsec (µs)

Packet size (bytes)	64	128	256	512	1024	1280	1518	
Port 1 to 2								FIFO
								LIFO (Calculated)[4]

[4] See the Latency discussion in section 3 of this guide for the LIFO calculation.
[5] Data not available yet.